UPSTREAM 3.0: CABLE'S RESPONSE TO WEB 2.0

Dr. Robert Howald and Phil Miguelez
Motorola Mobility

ABSTRACT

It has been a little over ten years since Upstream 1.0: Video Service Support migrated to Upstream 2.0: Internet Access. Continuing traffic growth coupled with consumer expectations mark the next phase of planning. In particular, today's web experience includes new behavioral and service type accelerants that continue to apply pressure on upstream capacity. On the behavioral side, the Web 2.0 paradigm is always-on, real-time, high-speed access that delivers the cloud of connectivity that consumers expect. The percentage of the population documenting their everyday lives minute-by-minute and tracking others doing the same has grown rapidly. On the services side, we have seen staggering growth in user-generated content (UGC) tied to the social networking phenomenon and in multi-media uploads for content sharing on sites such as YouTube. These image and video-centric experiences represent a substantial increase in bit volume. As a result, MSOs now must plan accordingly and deliver Upstream 3.0: Web 2.0.

In this paper, we explore the implications of Upstream 3.0 and options for successfully delivering on these requirements. The two components to consider are optimizing the efficiency of available spectrum, and preparing an upstream migration plan that can support an aggressive compound annual growth rate (CAGR). Optimizing available spectrum includes utilizing key DOCSIS® 3.0 tools such as wideband 64-QAM, S-CDMA's newest features, and pre-equalization. Positioning the network for anticipated CAGR involves recognizing its implications on the upstream lifespan and developing a strategy to address these implications. As CAGR marches steadily ahead, operators are at risk of running out of capacity within the business planning horizon. In this discussion, we will focus on the ability to harvest new DOCSIS® 3.0 capable spectrum up to 85 MHz and subsequently beyond. We will quantify the potential upstream capacity, illustrate its ability to deliver on long-term growth projections, and describe the implications in the plant and in the home. We will describe steps to ensure that these gains are fully realized, and to guarantee that potential additional expansion is cost-effectively achievable. Finally, we will describe recent HFC and in-home environment measurements that can guide operators in implementing the migration.

INTRODUCTION

With the introduction of high-speed data services over DOCSIS over a decade ago, signaling the end of the dial-up age, broadband users have easily found ways to consume the new bandwidth. At first, it was simply a faster web surfing experience. But shortly thereafter, the web experience itself was modified as creative entrepreneurs rapidly made use of the bandwidth in a surge of development to create the most compelling web pages, and large scale e-commerce opportunities emerged. While surfing is primarily a downstream experience, DOCSIS introduced a significant jump in upstream speeds as well, ushering in new behaviors and activities – in particular peer-to-peer file sharing. It was this peer-to-peer beast that drove the period of bandwidth growth beginning early in the last decade and lasting several years, in part as legal battles wound down.

As peer-to-peer wound down, however, important new social experiences were well underway, growing up from low-bandwidth but highly popular chat (AOL messenger, chat rooms) to media sharing experiences, online games, and user generated content (UGC) driven by the standard set by YouTube. The social and community networking experience was yet one more element created by these pioneering network speeds, but also by the responsiveness and what was gradually becoming ubiquitous access. Ubiquity began as finite set of anywhere's – coffee shops, airports, restaurants, bookstores, etc. outfitted with WiFi. Today, it extends to truly anywhere because of the universality of the smart device. Virtually every phone today qualifies as a Smartphone relative to years gone by. Along with the needs of businessmen and women, some Web 2.0 habits in particular (Facebook/Twitter Updates, Sports GameCenter/Fantasy, mobile photo uploads) have led to a 24/7 connected paradigm and ubiquitous, portable access. This always-on, always-connected, always immersed in social networks, experiences, and media swapping, is the driving force of the New Era of Internet usage, Web 2.0.

Web 2.0 is simply about how the modern Internet is used – augmenting classic search and consume browsing with usage behaviors around UGC publishing, social networking, media sharing, media-centric (Flickr) and community-oriented web services and applications (Groupon, ESPN Fantasy, PokerStars.net), etc. Real-time broadband two-way IP sessions are the basis of these activities. Along with a newly emerging interest in home automation and associated services, these activities are the new residential pressure on upstream bandwidth. DOCSIS is cable's IP connectivity system, and its capabilities are in large part responsible for there being the evolution we now call Web 2.0.

UPSTREAM GROWTH MANAGEMENT

Consumer Expectations and Trends

Authoring and publishing of content locally – UGC – in the form of pictures, music, and video, requires a robust upstream data service. DOCSIS enables a high-speed upstream; however, tremendous pressure will be placed on current DOCSIS capability to keep pace with continued growth. That the social networking aspect of web life has grown is not a secret, and that UGC has become a major part of it is also no surprise. However, Figure 1 puts into perspective the scale of that growth. It represents that "yellow" or "red" flag for operators trying to grasp what Web 2.0-type activities create for the HFC upstream that simple web page (browsing) requests did not. And, this is ahead of any surge of home automation services that may lie ahead.

Figure 1 compares various aggregates of professional media content as measured in Petabytes (10^15 bytes). The right-most bar represents an estimate of the largest single day of photo uploads to the cloud in the United States in 2008. This one *day* (July 4? Mother's Day?) of upload activity represented more than twice the entire *year* of original television content generated worldwide. The photo upload bar also towers over the universe of all-time movie releases (somewhat impressive) and music (more apples and oranges). Fortunately for cable, a significant amount of these uploads are synchronized with Smartphone universality and Web 2.0 behaviors, not simply family photo albums moving between homes and clouds. As such, a large burden for this upstream activity is also being placed on wireless networks. However, as wireless networks, and in particular 3G networks, become crowded with new traffic, offloading of the traffic to nearby WiFi infrastructure often will ultimately have some of it ending up on cable networks. Furthermore, MSOs are actively courting new WiFi traffic with distributed plant access points for broader wireless coverage, and to take advantage of the rapidly scaling count of portable IP devices with sophisticated processing capability.

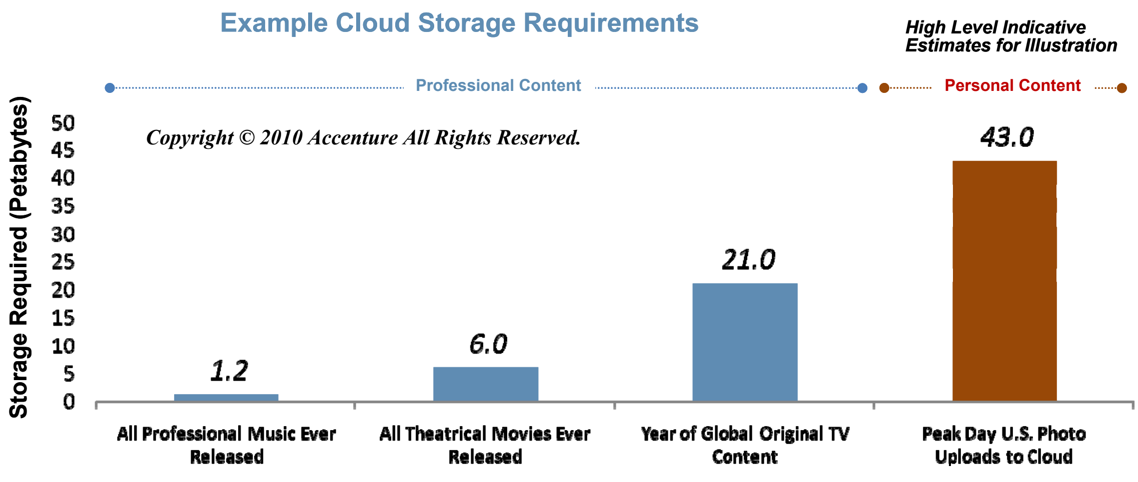

Figure 1 – Web 2.0: Upstream UGC Demand Two-Way Broadband

Tracking Bandwidth Growth

Traffic models based on Nielsen's Law, which uses a compound annual growth (CAGR) methodology, have been shown to represent well historical traffic growth trends. While actual traffic tends to grow and stagnate in a staggered fashion over the years, Nielsen's law has historically been a sound way to mathematically describe it, with the understanding that it works best as a long-term representation. As such, there is an element of "placing your bet" with regard to observing the historical trends and usage in the context of the new paradigm delivered by Web 2.0 experiences. Because of the short term fluctuations, there is a need to engineer ahead of the curve so as not to be waylaid by an unexpected step function in growth. There are several potential applications that we could ponder that could create such a step, but history has also shown that trying to guess the winners and the timing has been difficult. A new service or application often scales quickly and often catches us by surprise, creating "Napster" moments.

Figure 2 uses Nielsen's law under three different CAGR assumptions – 30%, 40% and 50%. It is Nielsen's law applied to a node or service group aggregate, under the assumption that average per-user increases are reflected similarly by the aggregate. The three trajectories are interrupted by two breakpoints over the next ten years. These represent node and/or service group splits – effectively 3 dB (best case) offsets, or a doubling of average bandwidth per home. Note that the 3 dB would be a step *straight* downward by 3 dB at implementation, so that by the time the next year comes around, some of that has been consumed, and the year-to-year step is less than 3 dB.

These trajectories are plotted against three different HFC upstream thresholds, using raw physical layer transport rate, as we will do throughout for consistency:

- 60 Mbps – Approximately two 64-QAM DOCSIS channels at 5.2 Msps
- 100 Mbps – Approximate available bit rate in 5-42 MHz with only A-TDMA
- 150 Mbps – Approximately a fully utilized 5-42 MHz using both A-TDMA and S-CDMA

With this information juxtaposed on Figure 2 with the CAGR trajectories, we can estimate at which point the various CAGRs will exhaust the available upstream throughput. The starting point of the trajectory – a key point from which all subsequent growth takes place – is an assumption of providing a maximum tier of 5 Mbps, while traffic engineering (TE) at 50:1 oversubscription on a 500 hhp node at 60% penetration. This nets out to providing 30 Mbps of total upstream to that service group, or one fully utilized wideband 64-QAM channel when hitting the 2% concurrent use metric at peak service hour.

The working assumption, then, can be looked as keeping pace with tier rate increases that follow demand-based CAGR, or managing the TE aggregate to keep pace with CAGR. It is a simple matter to shuffle the starting point on the chart up or down based on specific cases – not yet offering a 5 Mbps tier, averaging 5 Mbps with the economy tiers for a net average service rate, higher or lower rates of oversubscription and penetration, etc. The entire trajectory simply scales up or down by the same amount of dB offset. In addition, in keeping pace with user demand, more upstream comes with the price of delivering increases in downstream tiers. These increases must be supported by some consistent level of up/down asymmetry to be effective.

Finally, it is important to emphasize that it is a demand-based consumption analysis. Competitive market forces also have an important impact on the service rates offered.

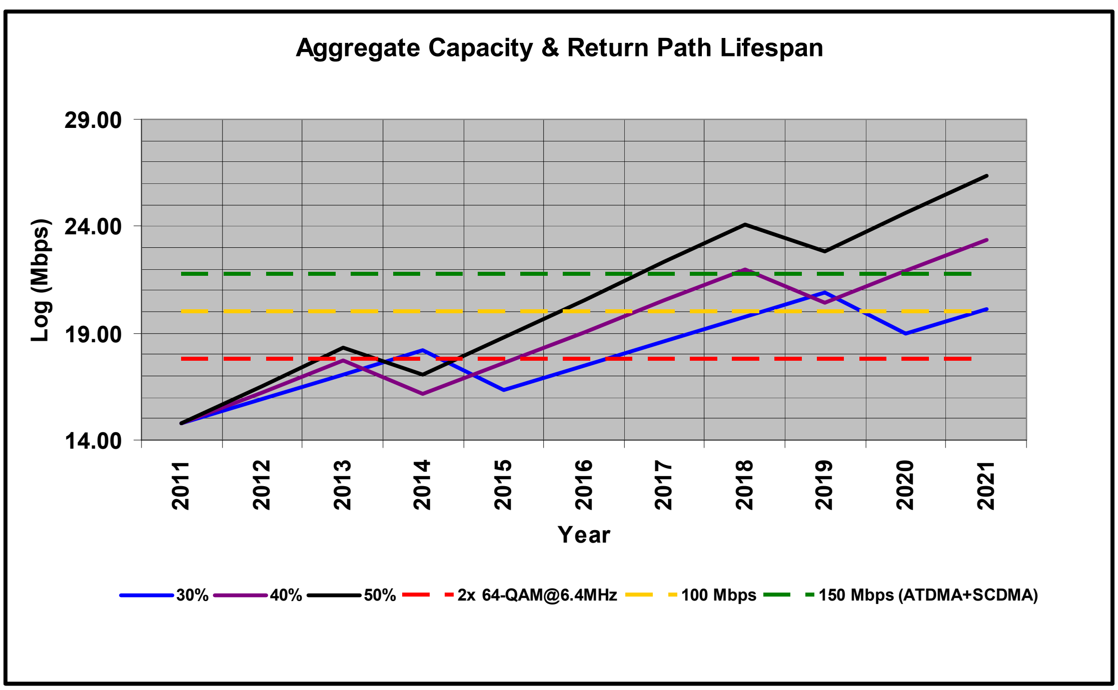

Figure 2 – Using CAGR to Project Upstream Growth Thresholds

Some key conclusions can be drawn from Figure 2. Clearly, a couple of 64-QAM DOCSIS channels get exhausted within a few years without a service group split. MSOs generally have recognized this and are planning a generation of splits or executing them on an as-need (traffic demand) basis already. More important, however, is to estimate when 5-42 MHz itself gets exhausted. While node splits are costly and intrusive, they are business-as-usual (BAU) type activities operationally. When demand or market needs push beyond what 5-42 MHz can provide, a more significant change must be considered. And, while further node splitting will indeed provide more average bandwidth, the maximum service rate limit also come into play as we look further into the future and project towards 100 Mbps upstream service groups.

For example, in an A-TDMA only case, a best case scenario may be providing a 100 Mbps peak service rate from the throughput available, with a collection of bonded 64-QAM channels providing the best effort service. One such service group could be enabled for upstreams well-behaved enough to support the bandwidth and modulation profile (64-QAM) required. Service rates significantly above 100 Mbps would be unattainable

Referring again to Figure 2, note that a single service group split gets us through 4-7 years of growth, based on CAGR assumptions shown, considering 100 Mbps as the 5-42 MHz throughput boundary. The former (4 yrs) represents a “should be planning what’s next” time frame, while the latter registers less urgency. Again, a second service group split would buy more time, but would not address peak service rate growth objectives. In particular, the ability to offer 100 Mbps to commercial customers would be limited in a 5-42 MHz only architecture, and shared with residential users. With diurnal variations of commercial and residential, this may be reasonable.

Note that through use of S-CDMA (green) and an assumption of relatively robust (40%) average growth, the upstream could last the decade. Recent data [1] suggest short term CAGRs have slowed, registering at about 25%. This represents, approximately, a three year traffic doubling cycle. Critical variables to the 40% result, however, include the aforementioned commercial services, and the lingering possibility of an aggressive CAGR (50%) bandwidth-buster type application, either of which could break this attractive conclusion. The 50% CAGR would pull the bandwidth exhaustion point in to 7 years. Finally, well before the 5-42 MHz is exhausted, planning will have had to begun, and in fact steps taken, to be prepared with new bandwidth. This is simply because the nature of any new steps involves some intrusive changes in the plant and home environments.

Figure 2 acts as a useful guide to managing growth versus time, and is relatively easy to extrapolate into other circumstances more specific to an operator or region.

<u>OPTIMIZING LEGACY 5-42 MHZ RETURN SPECTRUM</u>

Given the inevitability of a bandwidth bottleneck, with the possibility that it could be a near-term concern, it is clearly important to understand how to squeeze every last bit per second out of the 5-42 MHz upstream. Important underlying assumptions were made in the prior section in order to add up all of the possible return path bits per second and to determine thresholds to mitigate CAGR. In particular, we assumed 64-QAM upstream links at 5.2 Msps (6.4 MHz) were possible, that A-TDMA would limit the ability to extract throughput out of the 5-42 MHz return paths, and that S-CDMA opened up extra bandwidth where A-TDMA would not. Much has been written and discussed on these topics [2][3][4]. We summarize some of those findings here.

Optical Link Performance

Figure 3 singularly summarizes the HFC element of return path optics relative to their ability to support increasingly sophisticated modulation profiles. The return path optics, which typically sets the performance of the plant itself (Home and HE not included), are displayed from least capable to most capable, beginning with Isolated Fabry-Perot (IFPT, red trace), in each case using minimum specified performance. Note that yet older Fabry-Perot lasers exist, which are non-isolated. They have even lower performance. These are not shown here because they play no useful role in the enhancement of the return path for high-speed data services. They were deployed in their era for supporting interactive STB traffic, which is very narrowband and implements a very robust modulation profile.

From the IFPT, the additional traces identify minimum performance of the latest generation of FPs (EIFPT), two Distributed Feedback (DFB) lasers – 1 mw and 2 mw – and a digital return system based on 10-bits of transport (for this case only, measured).

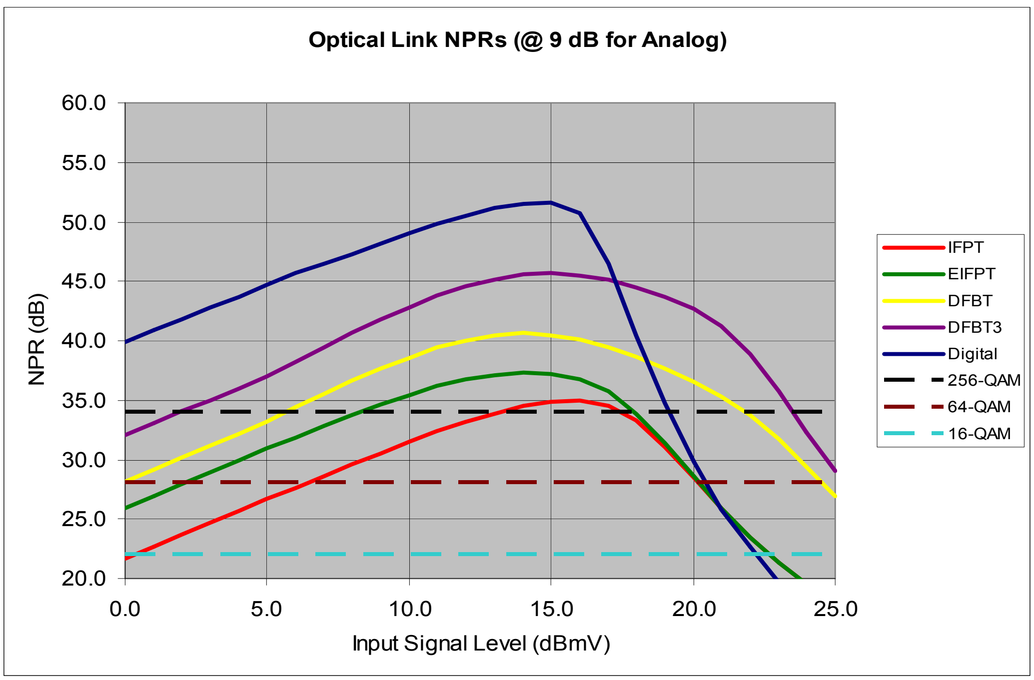

Figure 3 – Optical Link Noise Power Ratio vs Technology

The QAM thresholds shown represent 16-QAM, 64-QAM, and 256-QAM, the latter being non-DOCSIS. The values shown represent SNR values of 22 dB, 28 dB, and 34 dB, respectively. On the left-hand side of the NPR curve, away from the peak, NPR and SNR are one in the same. The area around the peak and to its left is the practical operating region of the curve. The concept of Noise Power Ratio as it applies to HFC returns is described in great detail in [5]. The basic NPR concept is that it represents the performance of the return link if it were fully loaded with digital channels, with the total power of those channels being the value on the x-axis.

The QAM thresholds chosen represent approximately 1e-8 *Bit* Error Rate values. From this, one can consider that there is forward error correction (FEC) coding gain that means several dB *less* is acceptable as a performance threshold, as is done for the DOCSIS downstream (23.5 dB/30 dB for 64-QAM/256-QAM post-FEC). Conversely, one can reserve coding gain for the unknown that often is symptomatic of the return channel, and add several dBs above the shown thresholds to identify what is a comfortable operating margin. Those policies vary across operations and system architects, so we use a value that comes with no gray areas to consider.

For reference, Motorola has taken the additional step of correlating measured *codeword* error rates (i.e. with FEC) to packet error rates (PER) – that which would reflect end user experience. The results indicate that an uncorrected *codeword* error rate of 1e-3 approximates the threshold beyond which a steady increase in packet errors ensues beyond an acceptable level, as shown in Figure 4. Many bit errors are required to make a codeword not correctable, however. Suffice to say that a 1e-8 BER would be error-free post-FEC. So, in that sense, use of the above SNR thresholds is conservative relative to packet loss. The interested reader is referred to [2] for more detail.

Referring back to Figure 3, we quantify in Table 1 the operating dynamic range for the two higher order modulation profiles provided by each of the return path optical technologies shown.

Packet Error Rate vs Uncorrectable Codewords

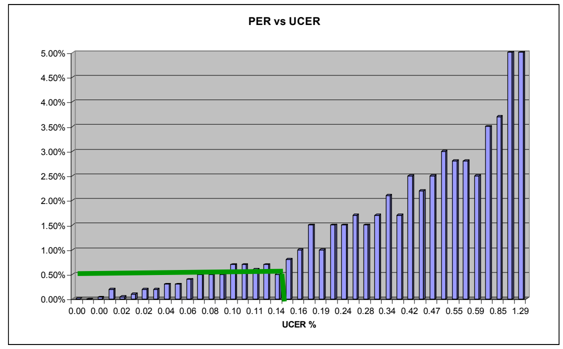

Figure 4 – Codeword Error-Packet Error Correlation

Table 1 – M-QAM vs Dynamic Range

	Dynamic Range	
Return Optics	64-QAM	256-QAM
IFPT	8	2
EIFPT	11	5
DFB (1mw)	15	9
DFB (2mw)	19	13
Digital (10-bit)	25	19

Several conclusions can be drawn from Figure 3 and Table 1:

- First, Fabry-Perot (FP) lasers start to lose comfortable margin levels when 64-QAM is deployed. Moving from 16-QAM to 64-QAM is at minimum 6 dB of lost margin. Link performance shown is further compromised by other built-in system noise, such as deep cascades and receiver noise figures, further degrading margin. We will quantify some of this loss in a subsequent section.

- 64-QAM can be supported over modern FP lasers, but the lost margin is more likely to make deployment challenging. It could result, for example, in inconsistent performance and/or maintenance due to impairments that do not bother 16-QAM. As such, getting started with existing FPs is quite possible and allows a gradual migration to DFBs to take place, but there must this migration must take place to fully maximize throughput with robustness.

- DFB optics and digital return optics comfortably support DOCSIS 64-QAM performance. Furthermore, these technologies offer the potential to increase the modulation order to yet more bandwidth efficiency using 256-QAM. In the case of the 2 mw DFB and 10-bit digital returns, double digit (i.e. comfortable) margin exists for supporting the additional 33% new bandwidth efficiency of 256-QAM over 64-QAM.

The moral of the optical link part of the "optimizing" story is simple – migrate towards DFB optics or digital return to fully utilize the 5-42 MHz return. Which direction to go includes various other factors that weight the pros and cons of analog and digital [6], in particular the capability to continually support more new RF spectrum to be discussed herein.

S-CDMA

It has long been understood that the low end of the return band is a messy place to live, fraught with short wave radio interference and home-induced impulsive noise from a variety of common sources. As such, most MSOs write-off the region of spectrum below 15 MHz for DOCSIS services, perhaps even below 20 MHz. Figure 5 shows a classic example of the kind of muck a signal may find itself suffering through in the lower half of the spectrum. In this case, the burst is wideband enough to intrude on a portion of the DOCSIS channel in a way that impacts short term Modulation Error Ratio (MER). Resulting MER behavior from such an event is shown in Figure 6.

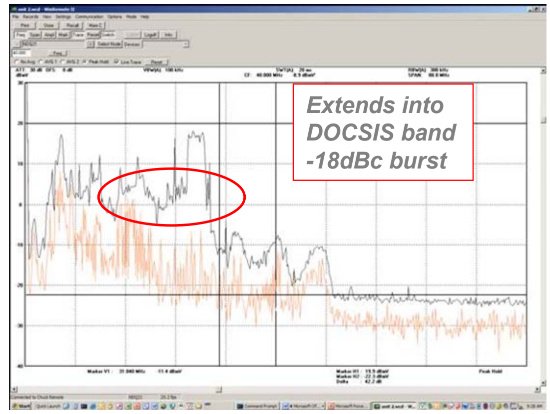

Figure 5 – Spectrum Capture at HE of an Impulse Noise Burst

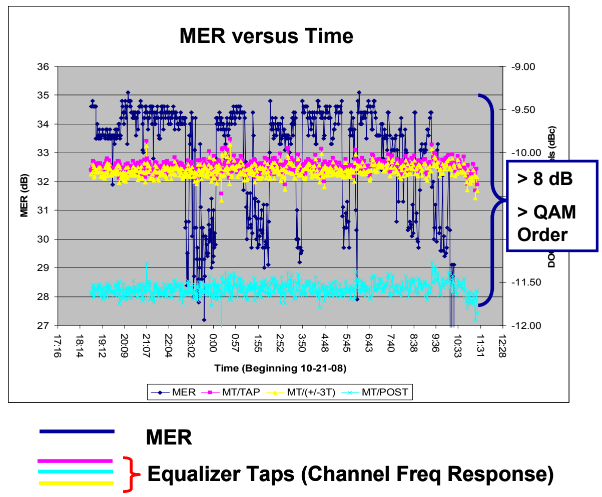

Figure 6 – MER Behavior Under Severe Impulse Noise

While inherently unfriendly territory, DOCSIS 2.0 put in place the mechanism to deal with this portion of the band. DOCSIS 2.0 includes the requirement to support S-CDMA. Figure 7 shows a recommended channel distribution [4] using S-CDMA to squeeze every last bit out of 5-42 MHz. In this case, S-CDMA offers 33% more capacity. Under complete avoidance of the low end of the band, up to 50% more throughput can be added compared to today's usage. We will utilize these S-CDMA conclusions to determine its role in extending the lifespan of the return path for the legacy 5-42 MHz bandwidth, as well as for subsequent new bandwidth growth.

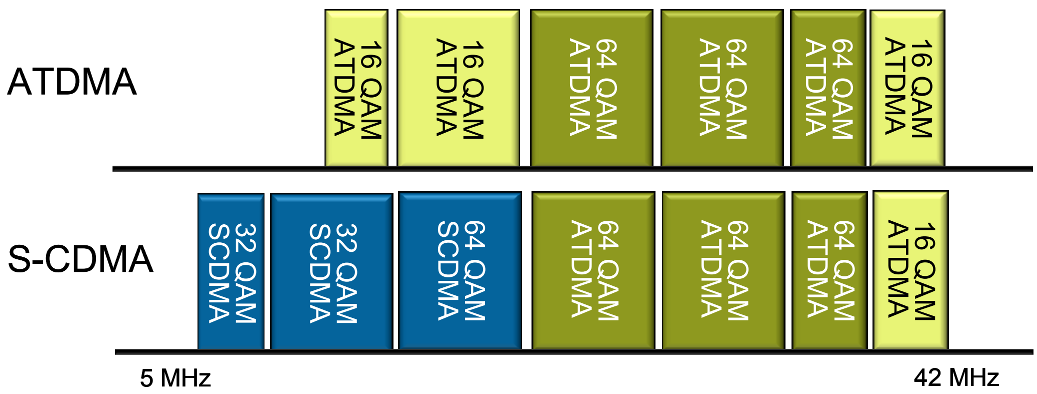

Figure 7 – Maximizing Available Throughput Using S-CDMA

<u>Capacity and Lifespan</u>

In Figure 2, we showed upstream CAGRs of 30%, 40% and 50%, and showed how those rates crossed various thresholds of available throughput in the 5-42 MHz band. We pointed out the critical nature of the CAGR variable itself – something anyone that has an interest-bearing savings account (is there such a thing anymore?) or had a securities account in the 1990's understands well. Dollars or Megabits, compounding math works the same way.

In Figure 8, we have displayed the same information in a different fashion, allowing us to understand the sensitivity of the exhaustion of the 5-42 MHz HFC plant relative to the CAGR assumption. The same finite set of threshold conditions are displayed, and we plot the curves from the same starting point as discussed for Figure 2 (5 Mbps peak, 50:1, 500 hhp@60%). In Figure 8, service group splits are instead represented by dashed traces for the 100 Mbps and 150 Mbps cases (only, for clarity).

The two crosshairs on the figure are positioned to help understand the interpretation between Figure 2 and 8. For example, observe the trajectory of the 50% CAGR and note the point at which it exhausts a 100 Mbps maximum throughput channel in Figure 2. This occurs slightly more than 4.5 years into the future. We can see this same point represented by the leftmost crosshair on the 100 Mbps (pink dashed) curve in Figure 8. Similarly, we can observe in Figure 2 that the 40% CAGR trajectory crosses the 150 Mbps threshold in seven years. This matches the crosshair marking the 150 Mbps threshold curve (dashed yellow) in Figure 8. In both cases, in Figure 2, these trajectories also crossed the thresholds mentioned after undergoing a service group split, thus the correlation to the dashed curves in Figure 8.

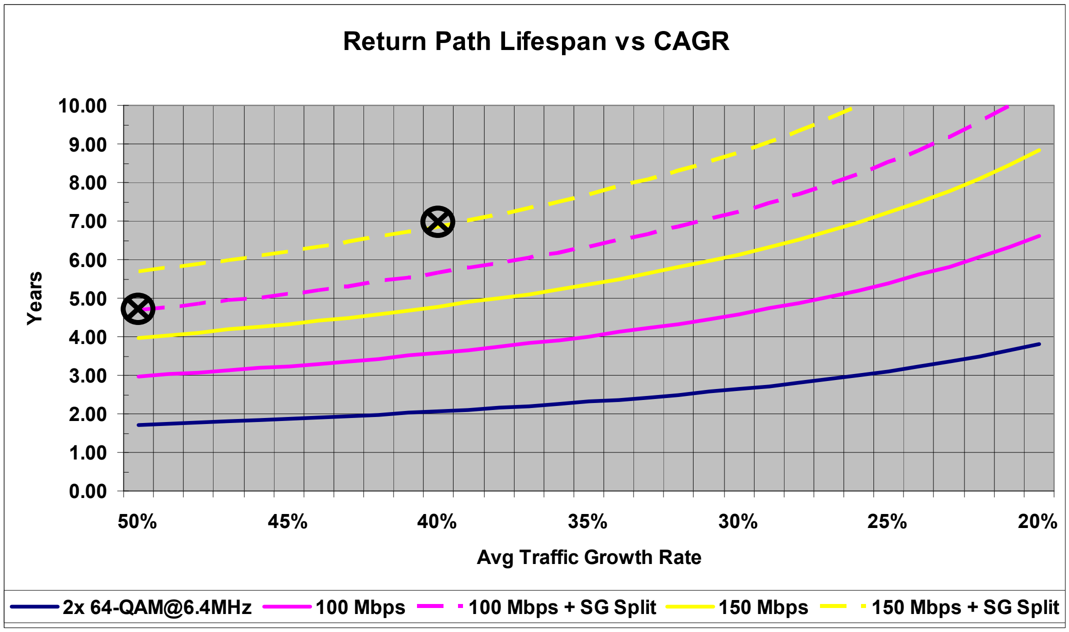

Figure 8 – Lifespan of 5-42 MHz vs CAGR

Figure 8 can be used as an excellent guide to the state of urgency (or not) of an operator's upstream situation based on historical and projected CAGRs. Its usefulness stems from the granularity provided by the perspective of Figure 8 when considering the sensitivity to the variable of most impact (CAGR).

Some of the core conclusions from Figure 2 can be drawn as well from Figure 8 under the given introductory assumption. A near-term need to add DOCSIS channels upstream will continue. At least one service group split is most assuredly on the horizon to provide support for average bandwidth growth, even under relatively modest CAGRs of 25-30%. This is the case in particular if operators continue to avoid the low end of the spectrum where S-CDMA can buy additional time.

As with Figure 2, Figure 8 quantifies average bandwidth growth from a demand standpoint. The same principles discussed around supporting increasingly high residential peak rates, 100 Mbps commercial service opportunities, and higher rates due to "market push," still apply.

DELIVERING NEW DOCSIS CAPACITY

Benefits of 85 MHz Return (N-Split)

An unavoidable conclusion from Figure 8 is that the end is in sight for upstream capacity. It may not be near (or it may be), but it is certainly to the point where planning for what comes next is prudent. For practical reasons, there is no intent to remake the 5-42 MHz modem technology to squeeze out any latent Mbps that may be possible given that DOCSIS is 10 years old, the RF section of the HFC plant has gotten shorter, and serving group sizes have gotten smaller. In fact, we have begun to take advantage of some of these HFC performance benefits within DOCSIS today by turning on 64-QAM modulation, which is a step closer to optimum use of the upstream channel compared to 16-QAM. And, any new modem technology developments that take place in the future for HFC and take advantage of the latest developments in communications technology and information theory to close the gap between current performance and theoretical capacity can always call out frequency range support for the legacy DOCSIS bands.

Table 2 illustrates the available DOCSIS transport rate for various split architectures, and the theoretically available capacity at the DOCSIS-specified minimum of 25 dB. While it is impractical to achieve theoretical capacity, the gap has indeed closed over time between practice and theory.

Table 2 – Bandwidth, DOCSIS, and Theory @25 dB SNR

Maximum Capacity for Each Bandwidth

Return Bandwidth	DOCSIS	Maximum Capacity
5-42 MHz	150 Mbps	300 Mbps
5-65 MHz	270 Mbps	500 Mbps
5-85 MHz	360 Mbps	650 Mbps
5-200 MHz	900 Mbps	1.6 Gbps

Note once more we using transport rate as the basis for all number unless otherwise mentioned. Actual user throughput upstream is highly dependent on upper layer (MAC) parameters, scheduling, packet sizes, and burst overhead.

Working down Table 2 from top to bottom, one obvious place to look for new capacity is simply new bandwidth. One straightforward way to exploit new bandwidth and remain compatible with DOCSIS is use of the 85 MHz return band, referred to as the N-Split. The limit set at 85 MHz was wisely chosen to maximize clean low band return without impinging on the FM band that could render some of the band difficult and unpredictable to use.

The advantages of considering expansion to N-Split are numerous. The primary benefits are listed below, and we will quantify several of them in subsequent sections.

- N-Split is supported by DOCSIS 3.0 for cable modems and CMTS
- Existing silicon is already capable
- Very clean new spectrum 42-85 MHz offers the potential for higher order modulations
- Legacy STB out-of-band signals (tunable across 70-130 MHz) are supported
- Entails minimal encroachment into the downstream band
- Has similar cable loss versus frequency properties as legacy band
- Multiple 100 Mbps tier serving groups are possible
- Support for traffic growth lasts through the decade under aggressive CAGR assumptions
- Architecture remains a diplex-based frequency domain duplex (FDD)
- Systems are possible with analog or digital return technology; analog already supports

We will compare additional properties of the 85 MHz approach to alternatives for supporting Gbps capability in later sections.

<u>New Return Spectrum: HFC Performance</u>

The opportunity to light up new, clean spectrum has the advantage of elimination of the impairments commonly seen in the lower half of the 5-42 MHz spectrum and shown in Figure 5. However, this added bandwidth does come at the expense of increased load to the return laser, and its fixed available allocation of total power. With the available bandwidth more than doubling, this nets out to 3.3 dB of SNR loss to the signals sharing the load over the fixed noise spectral density of the link. This loss is shown in Figure 9.

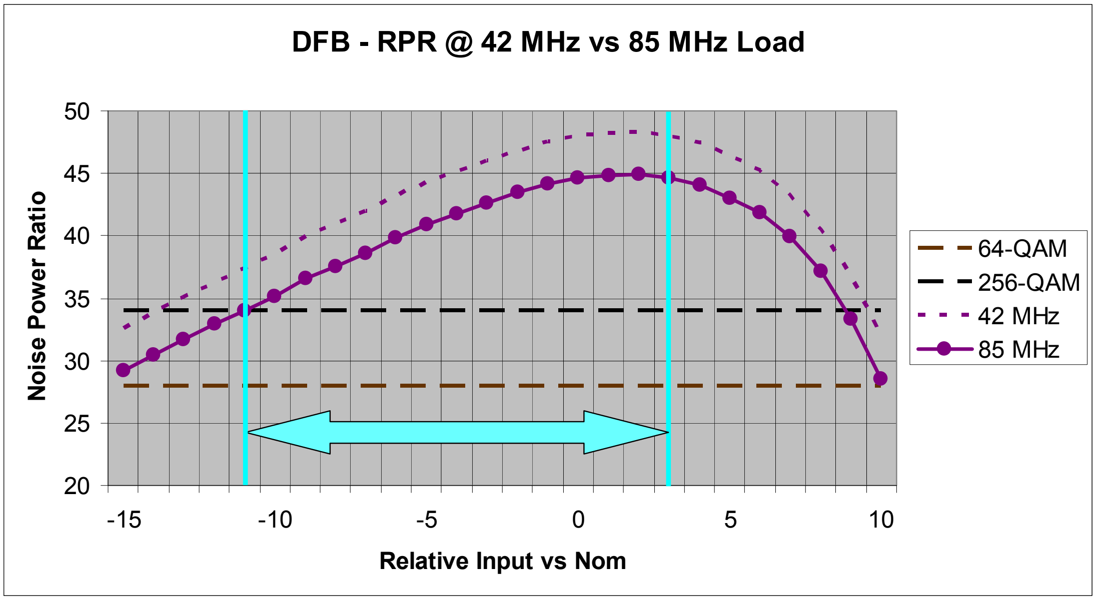

Figure 9 – Bandwidth Loading Effect (2 mw DFB Link @ 20 km)

Despite this loss in available SNR, it is apparent from Figure 9 that high quality return path optics – in this case, standard 2 mw, 1310 nm DFB lasers – have significant margin to support both 64-QAM and 256-QAM over typical HFC link lengths. In fact, we would estimate from Figure 9 that the HFC part alone has 13-14 dB of dynamic range over which 256-QAM could be supported in this scenario, quite operationally practical. As we shall see, this does not mean it is necessarily simple to flip the switch to 256-QAM. The HFC optical link represents but one component of the system, although it is the dominant contributor to performance of the plant itself.

Figure 10 shows a snapshot of a recent trial of an N-Split architecture, where the upper half of the band was used to support 256-QAM channels. In this case, supporting this modulation profile while co-existing with a maximum legacy 5-42 MHz load based on 64-QAM, was under evaluation. A mid-band test channel was left unoccupied for monitoring the most probable location of maximum distortion build-up as dynamic range was exercised, as in an NPR test.

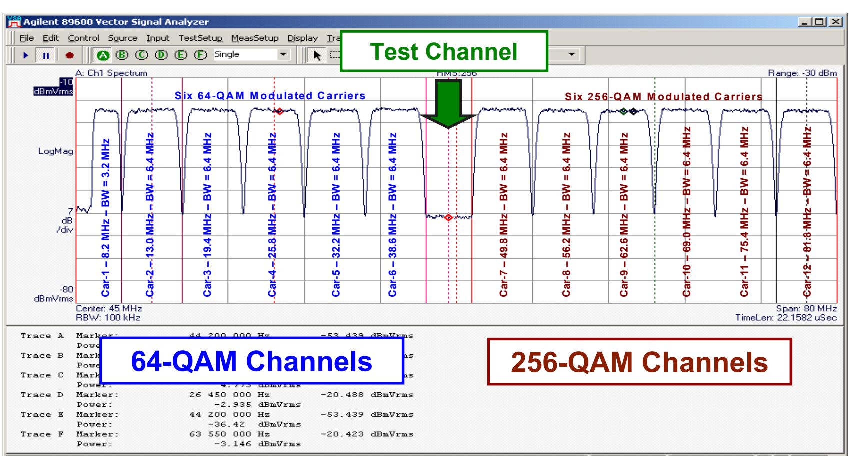

Figure 10 – 64-QAM and 256-QAM Channels Over N-Split

Evident from Figure 10 is the high available SNR delivered by the HFC link using existing analog DFB return optics at nominal input drive. The available SNR as measured at the input to the CMTS receiver of about 45 dB is consistent with the performance shown in Figure 9. In the case of Figure 10, the complete upstream link included a three amplifier RF cascade (i.e. an N+3 system). This performance is also consistent with support for higher order modulation profiles, such as 128-QAM and the 256-QAM example shown in Figure 9. Each represents a further increment towards closing the gap between theoretical capacity and practical data rates.

Considering the 5-85 MHz line in Table 2 and Figure 10, we can estimate that the gap between actual and theory is closed from 55% of capacity to 73% of capacity if migration across the band to 256-QAM were implemented. Of course, and as the 5-42 MHz line in Table 2 implies, extracting more out of the lowest end of the band is difficult. S-CDMA obviously helps considerably, as shown in Figure 7. However, with ideal channel capacity itself being a theoretical additive white Gaussian noise (AWGN) construct, it is to be expected that in the region where the noise is no longer AWGN-dominated that closing the gap would be more difficult. S-CDMA was leveraged at the low end of the band for the 85 MHz results shown above. Both with and without S-CDMA cases are quantified for N-Split analysis in subsequent discussion

Clearly, an inescapable conclusion from Figure 9 and 10 is that high performance analog optics – today's vintage of DFB-RPR links – are proven capable of supporting N-Split return links with higher order modulations for next generation upstreams.

New Return Spectrum: System Performance

As indicated, the HFC link is very capable of supporting higher order modulations over wider bandwidths than currently deployed. However, early generation CMTS equipment was designed to support 16-QAM as the maximum modulation profile requirement. Most vendors took this guidance and provided enough margin in their systems to enable 64-QAM, which was embraced in DOCSIS 2.0. However, enabling 256-QAM (an additional 12 dB of performance, minimum, over 16-QAM), was not cost effective to consider in early stages of DOCSIS deployment. This is not the case, however, in some newer receivers.

Given DOCSIS 2.0 requirements to support 64-QAM, as well as the requirement to support S-CDMA, which entails another degree of fidelity associated with its synchronous operation, new receivers coming to market, such as Motorola's RX48, have increased margin and dynamic range. Because of this, 256-QAM can now be comfortably supported. DOCSIS does not yet call out 256-QAM. However, much of the existing silicon base already supports this mode using the basic physical layer architecture of DOCSIS 3.0. Figure 11 quantifies the performance of the end-to-end system.

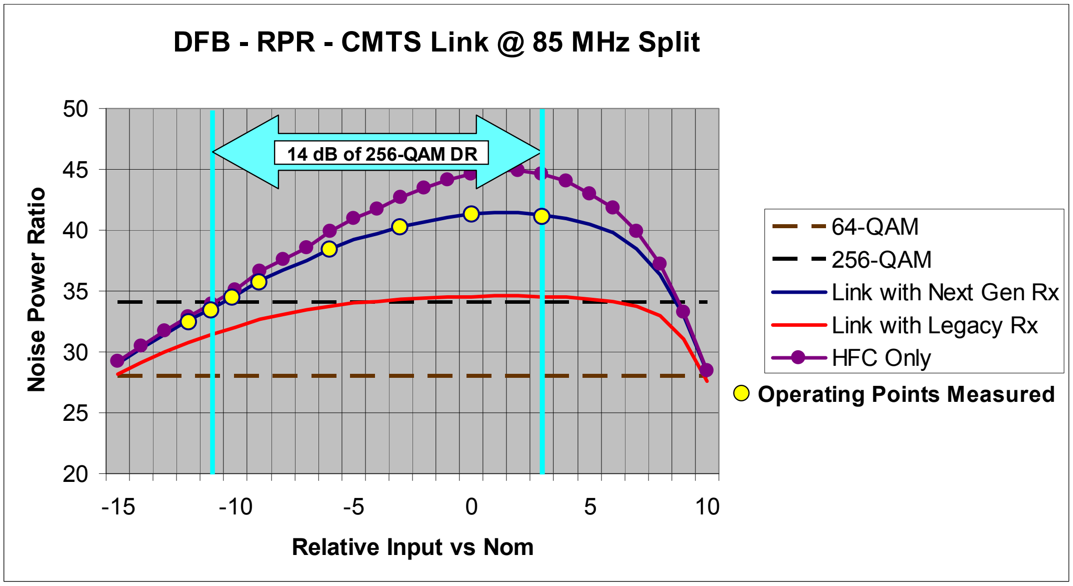

Figure 11 – 256-QAM Dynamic Range Over N-Split

Several key points can be taken from Figure 11. First, DFB HFC optics plus today's CMTS receivers comfortably support 64-QAM with sufficient, practical, operating dynamic range. This lesson is being proven everywhere DOCSIS 3.0 is being deployed. As previously described, in some cases newer, high quality FP lasers such as shown in Figure 3 can support 64-QAM as well. However, moving forward with more channels loaded and higher bandwidth, it is recommended that these eventually be replaced with DFBs. Nonetheless, it is comforting to know, given the magnitude of this task, that newer FPs can get 64-QAM started while the large task of exchanging lasers methodically takes place.

A second key point is that, using new CMTS receivers with extended dynamic range, such as the RX48 previously described, sufficient performance exists for 256-QAM to be practical. In fact, it is supported with nearly the same dynamic range that existing receivers provide for 64-QAM. Finally, comparing the HFC (purple) NPR trace to the HFC+CMTS (blue) trace, it is apparent also how little loss of NPR is incurred by new high fidelity CMTS receivers. Note that the yellow marked points on the composite NPR trace of Figure 11 represent low packet loss, *measured* loading points during the testing phase.

To complement the above field trial work performed with Cox Communications [7], Motorola chose to optimize loading further to determine the maximum throughput supported under the same HFC optics and the new generation of CMTS receivers. Both S-CDMA and A-TDMA were utilized, using 12 carriers employing modulations from 32-QAM to 256-QAM across the band. As before, however, 256-QAM was implemented only in the band above 42 MHz (and as such, still leaving some possible additional bits per second on the table). The 32-QAM recognizes that at the low end of the return band, even S-CDMA is challenged to overcome the noise and interference that congregates in that area. Table 3 quantifies the results of this evaluation.

Table 3 – Fully Loaded N-Split, A-TDMA + S-CDMA

5 MHz to 85 MHz Throughput Performance

	Frequency	Bandwidth	Symbol Rate	Modulation	Bits/sym	Data - SR	MOD	FEC-T	FEC-K	DOCSIS OH	ETH TP	MOD-PRO#
Car-1	11.4	6.4	5.12	32	5	25.60	S-CDMA	4	232	0.8242	21.10	431
Car-2	17.8	6.4	5.12	64	6	30.72	S-CDMA	4	232	0.8236	25.30	432
Car-3	24.2	6.4	5.12	64	6	30.72	A-TDMA	12	232	0.8724	26.80	522
Car-4	30.6	6.4	5.12	128	7	35.84	A-TDMA	8	232	0.9040	32.40	523
Car-5	37.0	6.4	5.12	128	7	35.84	A-TDMA	12	232	0.8705	31.20	524
Car-6	43.4	6.4	5.12	256	8	40.96	A-TDMA	10	232	0.8887	36.40	525
Car-7	49.8	6.4	5.12	256	8	40.96	A-TDMA	10	232	0.8887	36.40	525
Car-8	56.2	6.4	5.12	256	8	40.96	A-TDMA	8	232	0.9058	37.10	526
Car-9	62.6	6.4	5.12	256	8	40.96	A-TDMA	8	232	0.9058	37.10	526
Car-10	69.0	6.4	5.12	256	8	40.96	A-TDMA	8	232	0.9058	37.10	526
Car-11	75.4	6.4	5.12	256	8	40.96	A-TDMA	8	232	0.9058	37.10	526
Car-12	81.8	6.4	5.12	256	8	40.96	A-TDMA	8	232	0.9058	37.10	526
						445.44					395.10	

Raw Data Rate

Ethernet Throughput

The results of Table 3 indicate a maximum of about 400 Mbps of *Ethernet* throughput under the packetized traffic conditions used. For an apples-to-apples perspective with the other scenarios herein, and staying consistent with the rest of the paper, we will recognize the transport data rate of 445 Mbps as the value for comparative analysis.

Capacity and Lifespan

Having roughly doubled the amount of available spectrum, but also proven 256-QAM capability and turned on S-CDMA, what kind of additional lifespan can we expect out of a fully utilized N-Split, or at least as much as it has been proven to support? Using Table 3 results along with the 5-42 MHz thresholds identified in Figure 8, we add the new capabilities of N-Split to the mix in Figures 12 and 13.

Beginning with Figure 12, the gap between the set of 5-42 MHz options and the maximized N-Split is readily apparent. Even for the case where S-CDMA is not turned on to exploit the lower half of the return band, which most MSOs have been hesitant to implement to date, the difference between 5-42 MHz and 5-85 MHz capability is noteworthy.

The gap identified by the red arrow shows the lifespan impact for the CAGR choice of 30%. This value is lower than what has been seen during past periods of high growth, and more aggressive than has been seen during other periods, including the past couple of years. It could be considered a reasonable average over the past 5-7 years, if not slightly on the aggressive side. If S-CDMA is fully leveraged in both cases (42 MHz and 85 MHz), then Figure 12 predicts over four additional years of growth is available. Without use of S-CDMA (dashed green for N-Split) for either case, that buyback expands to nearly 5.5 years. In both cases, the transition to N-Split pushes the lifespan of the return path architecture such that it supports a full decade of new growth – a very comfortable chunk of next generation network planning time.

Figure 13 postulates that, given the time frames we are discussing in Figure 12, it is likely and even planned in many cases that node and/or service group splitting will occur, allowing for an increase in average bandwidth per home for the same total capacity. Both cases – with and without a split, are shown for two threshold cases in Figure 13 – the 100 Mbps use of 5-42 MHz (no S-CDMA), and the N-Split, also with no S-CDMA.

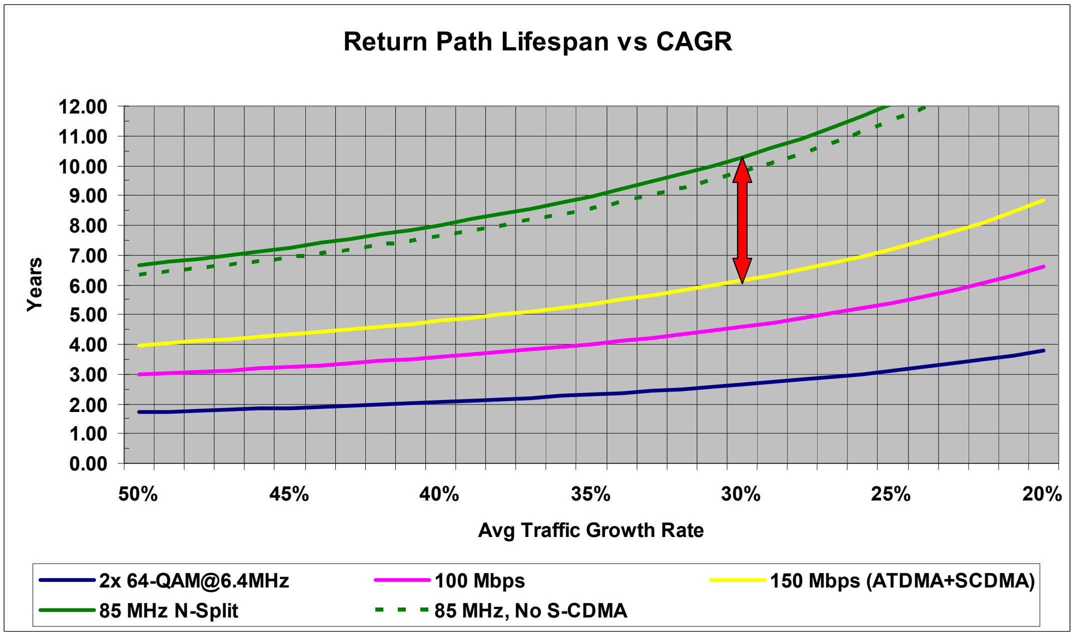

Figure 12 – N-Split Years of Growth vs. 5-42 MHz Use

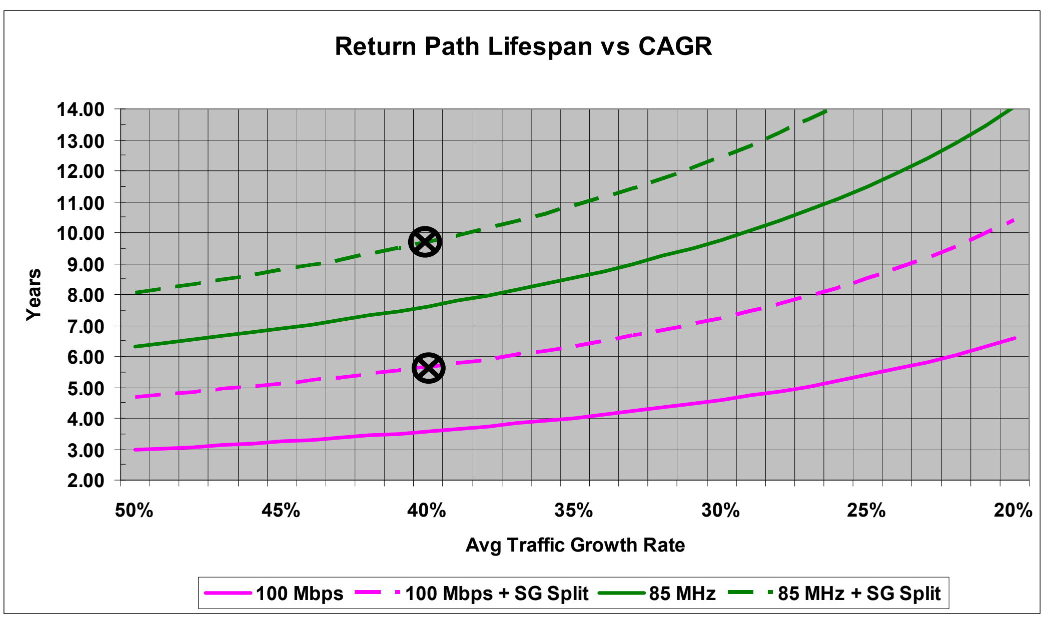

Figure 13 – N-Split Years of Growth vs. A-TDMA Only Plus Segmentation

Figure 12 showed the added lifespan for a CAGR of 30%. With a node split factored in within the next few years, as a prudent way to extend upstream support over 5-42 MHz, we can see in Figure 13 that the N-Split architecture now extends the life of the network to over 12 years of new growth at a 30% CAGR – beyond a reasonable period of planning the next implementation step given the speed at which technology change takes place. Because of this, when combined with a service group split, N-Split represents, for practical purposes of demand-based growth, a *long-term* solution.

Figure 13's bull-eyes highlight a more aggressive planning CAGR – in this case 40%. Again, under an N-Split upgrade, and without use of S-CDMA, almost a full decade of new growth is supported assuming a node split occurs. Obviously, with a second segmentation factored into the 10-yr span, more than a decade is then covered. Given the cyclical nature of plant investment and historical segmentation patterns, this may be simply a continuance of BAU bandwidth remediation. Alternatively, given the introduction of more cost-effective fiber architecture migrating into HFC, perhaps further HFC upgrades instead begin to give way to fiber-to-the-premises investment before the decade is out. An attractive feature of the N-Split is that it offers so much additional observation and planning time to consider the next phases of infrastructure investment.

The bull's-eyes themselves identify the basic mathematical relationship underlying compounding growth for a 40% CAGR. This value represents approximately a two-year traffic doubling period. Considering that we have used 100 Mbps to represent the A-TDMA only case within 5-42 MHz, and that the N-Split minus the S-CDMA contribution in Table 3 is almost 400 Mbps, we can back-of-the-envelope calculate the lifespan effect. The ~400 Mbps possible is 4x the 100 Mbps, so it is simply two traffic doubling periods. Since a doubling period at 40% CAGR is two years, two doubling period is four years. This is exactly the difference shown in Figure 13 as the lifespan bought moving from 5-42 MHz to N-Split at 40% CAGR.

<u>Legacy CPE Challenges</u>

The effects in the plant are easily understood for an N-Split migration. They are not necessarily operationally attractive, given that actives in the plant employ diplex filtering that must be changed to support 85 MHz. However, the steps to doing so are easily defined and represent no technological challenge. Instead, it represents primarily cost and logistics challenges. Managing the when and how in conjunction with fiber deep upgrades or node segmentations is an effective way to "kill two birds" and smooth the path for N-Split.

While plant migration is "blocking and tackling" steps, the implications within the home itself of modems transmitting in the new band is less well understood. Prior efforts have quantified the effect of 5-42 MHz modems on TV tuners when DOCSIS was first introduced [8], and recognized that, unsurprisingly, you could indeed overload a TV tuner even given an FDD architecture designed to keep forward and return apart from one another. With transmit frequencies now potentially in-band of tuners on CPE devices, it becomes yet more important to quantify the behavior of receiving STBs and TVs in order that new N-Split deployments do not disrupt current video services. As is the case in xDSL environments, it is anticipated that filtering in the home will be required in some cases to prevent degradation to video services. This would logically occur with a new modem deployment supporting N-Split, and would become a normal operation during such an install. However, the amount of filtering required is important to quantify, as well as the most sensitive frequencies for operational planning. An understanding of the nature of the observable impact for future customer support would also be valuable.

Figure 14 shows a block diagram of a test setup used to evaluate the sensitivity of different STB families and video loads to single and bonded channel transmit frequencies in the 42-85 MHz band.

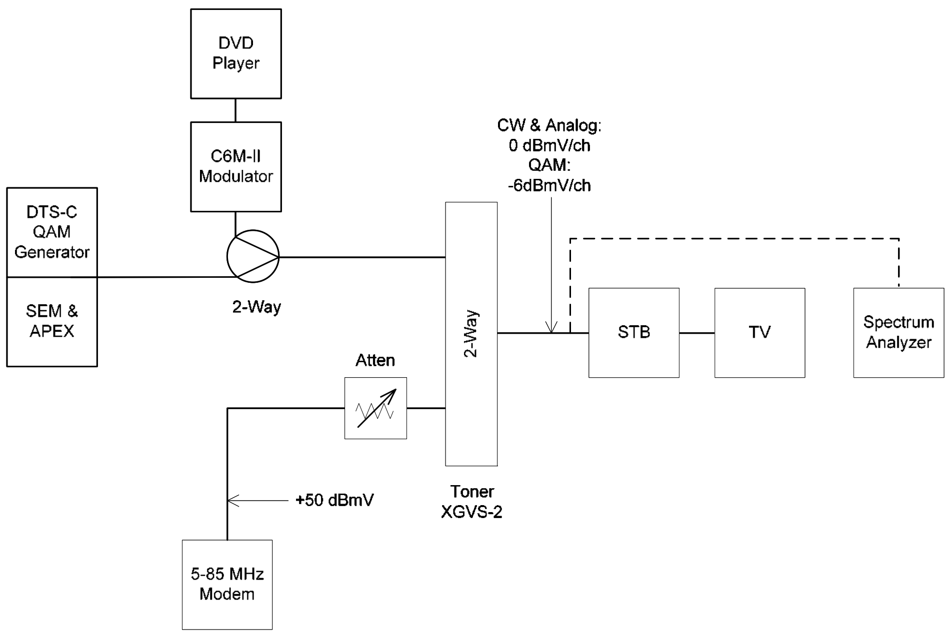

Figure 14 – CPE Testing for N-Split Home Environments

Figure 15 summarizes coarse results of the evaluation of video services that have analog, and those with only QAM. Based on observable pixilation or analog video picture distortion, it shows how much CM transmit power and what bands have the potential to cause interference concerns in the home, as a function of the RF isolation between the CM and the CPE. Of course, the amount of isolation is unpredictable in the home, and off-the-shelf retail splitters can have very poor port-to-port isolation. Figure 15 identifies (blue dashed lines) the range offered by splitters isolating CM and CPE of 10-30 dB, representative of a very poor or very good single splitter. Also identified is a typical range of CM transmit powers, up to the maximum allowed value (red dashed lines)

As to be expected, the STB is more robust against interfering signals when the channel load is QAM, allowing CPE signals as high as +30 dBc above the tuned video level in most cases. For analog services, the degradation threshold varied, but was on the order of 15 dB worse than in the QAM-only case. There is of course a frequency dependence between CM transmit frequency and CPE tuned frequency. There are many more permutations of this evaluation that will take place to comprehensively understand the impacts in the home, but these early results are illustrative of common tuner sensitivities. An important early conclusion, as can be seen in Figure 15, is that even under the most troublesome interference case of highest CM transmit power to most sensitive analog tuned input, the filter required to mitigate the interference requires a very reasonable 40 dB rejection value. This represents a relatively modest low pass filter design, capable of being easily manufactured at very low cost at these frequencies.

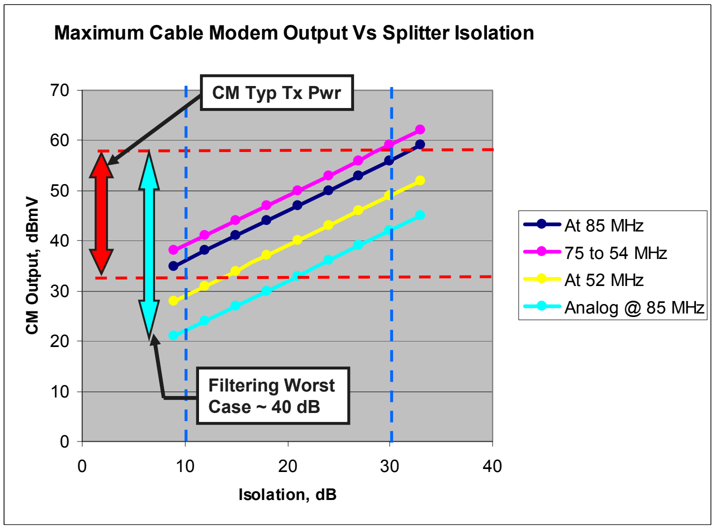

Figure 15 – CPE Testing for N-Split Home Environments

NON-DOCSIS SPECTRUM ENHANCEMENT

Gigabit Services

We discussed at length the ability of the N-Split architecture to support upstream traffic growth. We pointed out how, when coupled with BAU node splitting, how it should be considered a *long-term* traffic growth solution, not an upstream band-aid. The one feature that simply cannot be accomplished with the N-Split architecture is support of a full 1 Gbps of capacity, or naturally, of a Gbps tier rate. This is absolutely the case with DOCSIS use of the band (360 Mbps), and also the case even should theoretical capacity be achieved under DOCSIS SNR assumptions (650 Mbps). A theoretical 1 Gbps in the N-Split architecture alone would require a 38 dB return path SNR.

In Figure 11, we can see that with a new generation CMTS receiver and analog DFB optics, this 38 dB value is actually achieved over about 9 dB of dynamic range. In practice, however, that range essentially evaporates completely as theory gives way to actual implementation, and HFC variations eat away at what is left. However, it does point out that we are entering a new realm of possibilities on the return, where very high fidelity has never been the strong suit of the channel. Now, with 85 MHz of spectrum, modern HFC optics, and new CMTS receivers, many new dBs have been freed up that get us closer to capacity, and more importantly provide avenues that support new bits per second in the future.

Table 2 points out that a 1 Gbps threshold requires the split to move up to about the 200 MHz range. That bandwidth supports well over 1 Gbps of theoretical capacity, but under the reasonable assumption that DOCSIS remains in use in the 5-85 MHz band, the 85-200 MHz region can be exploited with something more aggressive that puts the total capacity above 1 Gbps. DOCSIS' maximum profile today (64-QAM@6.4 MHz) itself filling the band entirely out to 200 MHz falls short, although if 256-QAM were employed, even ignoring the low end of the legacy return, this would no longer be the case.

In the case of using split technologies (5-85 MHz of DOCSIS and 85-200 MHz of something else), a shortcoming that could come into play is the inability of that architecture to support a 1 Gbps *peak* service rate. It may be a complex endeavor in any case, but the complexity would be significantly magnified if any new technology developed to exploit 85-200 MHz needed to be integrated with the DOCSIS band to deliver 1 Gbps tier rates.

As was done in Figure 9, the additional bandwidth load must be quantified to understand exactly what a 200 MHz system might deliver. Figure 16 shows how the 200 MHz loaded performance compares to the 85 MHz and 42 MHz analysis previously discussed. The loss is again easily predictable, as simply the dB relationship among total bandwidths. And, similar to Figure 9, we can observe for this case that 10-11 dB of dynamic range for 256-QAM exists across the HFC optics – again an operationally practical amount of margin to accommodate alignment and plant behaviors.

Figure 17 is the analogous figure to Figure 11 for N-Split, showing, in this case, projected performance on a 200 MHz "high" split when factoring in an "equivalently performing" CMTS receiver (DOCSIS does not extend to 200 MHz). As would be expected, with the receiver performance equivalent to legacy CMTS receivers, inherently not equipped for 256-QAM, performance does not even breach the threshold. However, with a new generation of high fidelity receivers that achieves the noise performance of today's new cards, which already support 85 MHz, system analysis projects that there remains the full 10-11 dB of operational dynamic range to 256-QAM performance over a fully loaded 200 MHz return path.

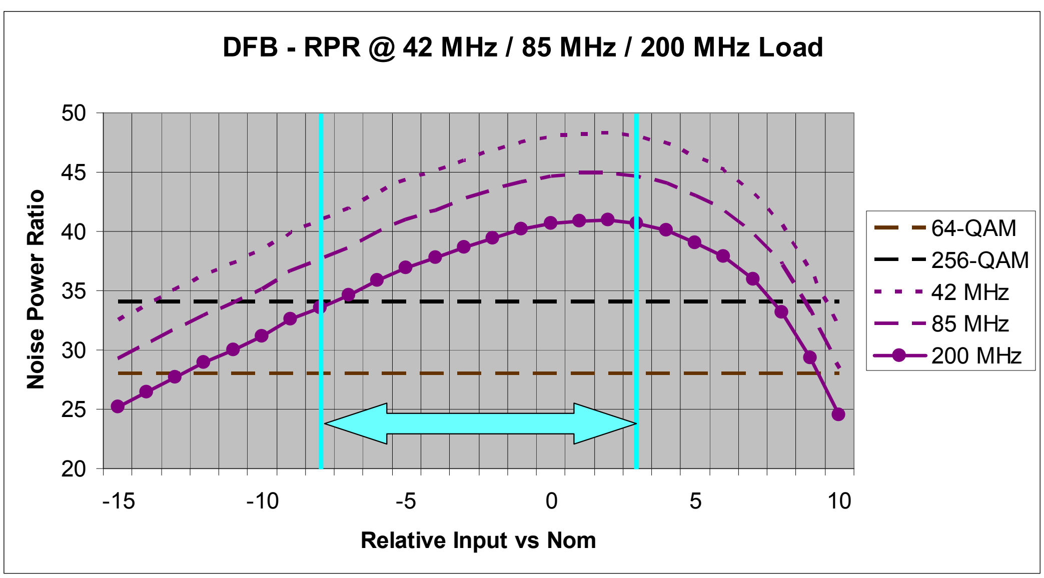

Figure 16 – Bandwidth Loading Effect, 200 MHz (2 mw DFB Link @ 20 km)

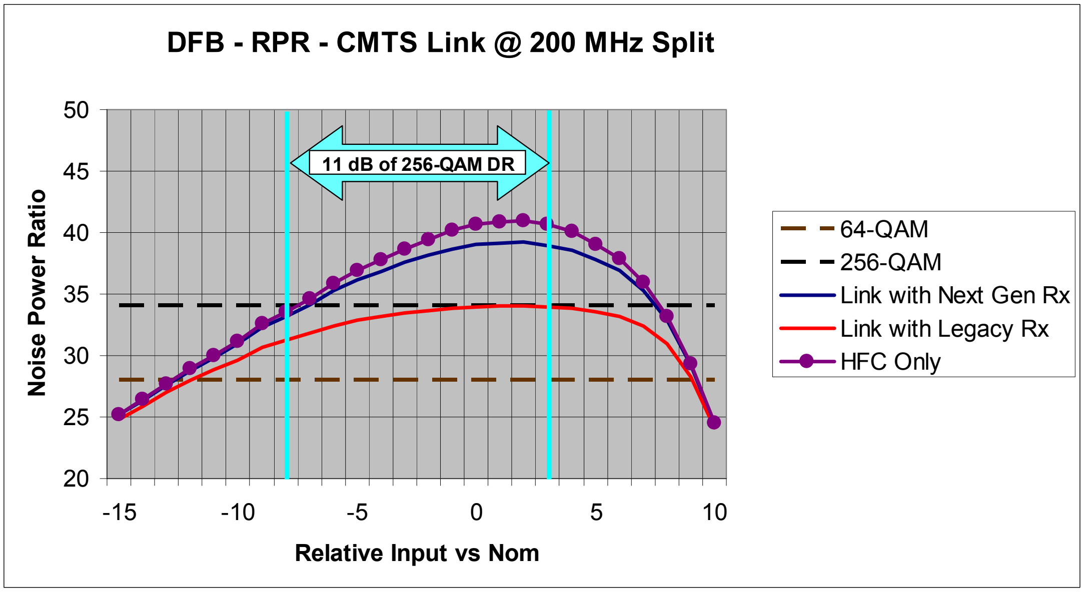

Figure 17 – Projected 256-QAM Dynamic Range Over 200 MHz Split

Alternative architectures exploiting coaxial bandwidth above 1 GHz have been around for many years. New iterations of these approaches could be leveraged to turn on currently unoccupied spectrum for adding upstream. The advantage of this approach is that in principle it does not interrupt legacy services, making a transitional path in theory non-intrusive to customers. Disadvantages include the need to work around legacy plant devices that are incapable of processing signals in this band, high losses at frequencies above 1 GHz, translating to significantly more power required from CPE devices, and new technology and operational hurdles that may arise in the new band. And, of course, the complete CPE itself becomes entirely new, or at a minimum requires the addition of block converters to support frequency translation.

On the passive network, coaxial cable and even some current 1 GHz taps are indeed capable of supporting useful bandwidth above 1 GHz [9]. However, the frequency dependence of cable loss quickly attenuates signals above 1 GHz when we consider it relative to the low band upstream. Given well-understood loss versus frequency relationships, for nominal trunk and drop spans of a passive segment, we can anticipate almost twice the loss (in dB) extending the return band to 200 MHz. However, above 1 GHz, the loss is increased by roughly a factor of five compared to legacy return for such a span. CPE devices must make up for that loss, and also must deliver additional total power associated with the wider bandwidth they would occupy to enable peak rates of a Gbps, relative to today's maximum of 6.4 MHz single channel power. Bonded channels as implemented today do not increase the transmit power.

Nonetheless, as technology continues to advance and HFC networks continue to get fiber deeper and RF shorter, the opportunities to exploit bandwidth above 1 GHz, whether for downstream or upstream, improves [9].

<u>Capacity and Lifespan</u>

So, what does all of this effort around making it to a Gbps of capacity and/or a Gbps of peak service rate gain for us relative to long term traffic growth? The answer to this question can be examined in Figure 18. It shows three threshold cases – 100 Mbps (A-TDMA only), N-Split (in this case, including use of S-CDMA), and 1 Gbps of capacity, however we manage to achieve it. Figure 18, as with any of the lifespan charts, does not care about implementation.

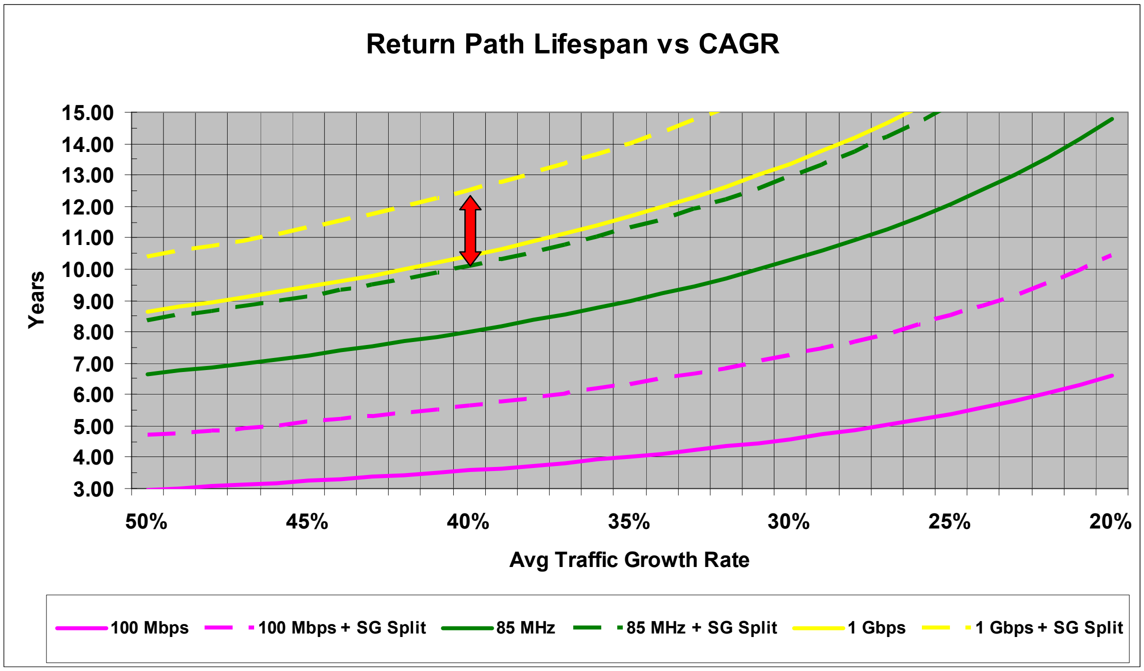

Figure 18 – Years of Growth: A-TDMA Only, N-Split, 200 MHz High Split

Zeroing in on the red arrow identifying the gap between N-Split and 1 Gbps at 40% CAGR, in each case with a node split assumed in the intervening years, we see that there exists about 2.5 years of additional growth. When we think of 1 Gbps, this intuitively seems odd. Why does migrating to N-Split buy a decade or more of traffic growth coverage, yet implementing a 1 Gbps system offers only a couple more years of survival on top of that decade? Unfortunately, this is simply how compounding works. It is up to our own judgment and historical experiences to consider how valid it is to be guided by the rules of CAGR and Nielsen, and if so what reasonable year-on-year (YOY) behavior assumption to assume. However, the mathematical facts of CAGR-based analysis are quite straightforward: with CAGR behavior, it takes many YOY periods to grow from, for example, 5 Mbps services today, consuming or engineered for perhaps tens of Mbps of average return capacity, up to the 445 Mbps we saw delivered by N-Split in Table 3. Once there, the subsequent annual steps sizes are now large.

As an example, if 20 Mbps of average capacity is needed to satisfy demand today, then traffic can double four times and not eclipse 445 Mbps. It eclipses 445 Mbps in the 5th traffic doubling period. For 40% CAGR (recall, two years doubling), that's a total of ten years. However, once the upstream capacity *has* consumed the N-Split, it is not very long via compounding before the 1 Gbps threshold is eclipsed. In fact, traffic can only double once and stay within this threshold. This is what Figure 18 is pointing out graphically. As such, relative to a solution that provides 1 Gbps, N-split gets us through 80% of that lifespan under the assumption of 40% CAGR and an intervening node split.

N-Split Plus

Based on analyzing the lifespan of the network to support traffic growth, the above result argues that N-Split is a reasonable next step that supports many, many years of new growth. It cannot achieve a Gbps of capacity or peak rate, at least as we can conceive of its use over HFC under reasonable system assumptions. However, it is not clear that Gbps service tiers will be required for residential upstream in a time frame reasonable for business planning purposes, certainly from a consumption demand perspective. Given the upstream lifespan offered by N-Split, evaluation can continue around what should follow, with technology changes in the intervening years likely to have a strong impact on the direction to head so many years down the road.

Nonetheless, while developing N-Split architectures, it seems prudent to enable a subsequent HFC step that can be implemented in a way that does not require another round of equipment visits to the physical plant. This can be used should Gbps capacity become critical as a marketing tool or for suddenly accelerated demand over a longer than normal period of time (i.e. a prolonged acceleration of CAGR). It is then also in place if it turns out that the best longer term evolution answer is that the HFC network lives on further, without a need to turnover technology such as to an FTTP architecture. If this is the case, and the decade plus of N-Split becomes exhausted, a natural "phase 2" RF step for upstream would be to increase the split once more. When considered as a possible phased solution, such an architecture would have selectable diplex options as part of the migration to N-Split. Multiple options could be made available under the control of the operator.

Another reason to consider a phased approach is for the practical reason of the legacy OOB downstream STB signals. These are tunable only to 130 MHz, so an upstream that engulfs this band obsoletes many existing STB. However, over the period of time that is bought by N-Split, these devices can be managed down, either through natural attrition or an active effort to reduce reliance on legacy OOB (vs DSG, for example), or via a more wholesale transition to a new generation of home gateways and/or DOCSIS-based IP video architecture.

Should 1 Gbps become part of the playbook for residential upstream, the approach to delivering this by extending the return band on the low side has many technical advantages over an approach using above 1 GHz. We have mentioned some of them already:

- Much better loss properties supporting more cost effective CPE
- No technology or plant hurdles such as housing, connectorization, bypassing
- Less guard band lost spectrum – no "triplex" guard band
- No bookending of downstream

Forward Relief

Moving to N-Split adds 43 MHz of return bandwidth, but does so at the expense of forward bandwidth. When factoring in the new guard band, possibly nine or ten 6 MHz slots in the traditional analog band are eliminated. Mathematically, converting these channels to digital allows them to all fit into one slot. As such, as analog reclamation continues, this forward loss does not represent a significant capacity concern. The primary operational concern is that the nature of the channels in this region often represent a basic service tier, and therefore cannot simply be transitioned into the digital tier and off of the analog tier, as perhaps some of the longer tail of the analog service could. Instead, some channel re-mapping and/or more aggressive deployment of digital adaptors would be required. In any case, given the powerful set of tools available to provide downstream capacity, 85 MHz presents no significant imposition on the forward bandwidth in terms of capacity loss.

In the case of a 200 MHz extension, however, this is not necessarily the case. This can easily shown to be so, for example, in 750 MHz systems that are trying to accommodate aggressive CAGR, such as to support OTT growth, where extensive analog exists and may continue. The issue is magnified further when considering new trick play video services and alternate screens that result in more unicast delivery, and when considering the addition of more HD, 3D, and possibly even higher resolution services.

As previously indicated, in the case of 1 GHz, there is significant "free" bandwidth available above the specified 1 GHz value. Figure 19 shows the frequency response on the "through" port of a particular 1 GHz tap – the port that would be in series with other taps on the way to a connected home. The response on the tapped port also has essentially parasitic, low-loss properties over the first 200 MHz above 1 GHz. Though not as perfectly flat, it creates no significant burden to RF signals in the band, and in particular when considering a new generation of modem technology, such as multi-carrier [9]. The same is the case for some families of 750 MHz taps (available bandwidth exists above 750 MHz) and 870 MHz taps (available bandwidth exists above 870 MHz).

The amount of useful bandwidth and loss properties are vendor dependent, but MSOs already often use slots above these limits. Conveniently, as Figure 19 shows, the amount of available new bandwidth simply trickling over the top of the band is virtually the same the amount of bandwidth that would be removed from the forward by 200 MHz systems, when considering guard bands for that extension. Note, however, that there is no forward/reverse guard band involved here without an upstream system contending for spectrum. Also, this "replacement" bandwidth amount provides adequate spectrum to facilitate downstream Gbps services. The ability to fully exploit this bandwidth in the passive plant obviously depends heavily on the band coverage of the actives themselves and the depth of the cascade. Clearly, this is where shortening cascades and "N+small" continue to payoff for HFC evolution.

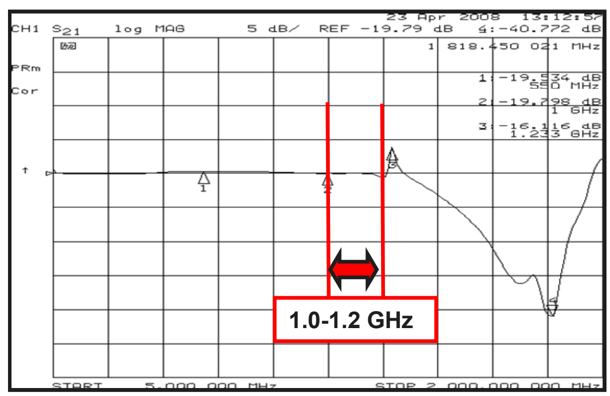

Figure 19 - Above 1 GHz Frequency Response of the Thru Port of a 1 GHz Tap

There are three other compelling advantages to considering use of the band over the end of the defined tap bandwidth for forward services, as opposed to reverse. First, the forward path is inherently designed for high fidelity in support of analog video. As we know, the reverse path was never architected with high fidelity in mind. Over time, technology has been introduced to enable a high-speed data channel, but the low noise and high linearity architected into the forward path is orders of magnitude above the return path. This difference translates to a much more straightforward exploitation of bandwidth with high performance on the downstream.

Second, the forward path levels are designed for RF path losses out to 1 GHz. Because of this, the parasitic losses above 1 GHz of the coax, and the minimal additional attenuation, are not a stretch to achieve when extending the forward path. It is an entirely different case in the return, where the architecture has relied on the low loss end of the band, which increases only modestly as it is extended to 85 MHz or even 200 MHz.

The third point, related to the first two, addresses the issue of CPE transmit power to overcome these high losses. Forward path RF systems, being design for similar losses, have seen investment in broadband RF hybrids drive higher and higher levels over increasing forward bandwidths, still based on supporting a full analog and digital multiplex. As a result, the output levels of these hybrids and nonlinear characteristics have continued to improve. However, investment in these premium devices for the forward path is spread over the number of homes serviced by the actives. More broadly, investment in the number of premium RF hybrids from the node to the final amplifier is shared by the number of homes passed by the node. In the reverse path, each home needs a high power, linear transmitter (though less than an octave) in this higher frequency band.

Figure 20 quantifies what is ideally available exploiting the frequency response above 1 GHz for forward band purposes based on the passive segment only, or representing effectively and N+0 situation.

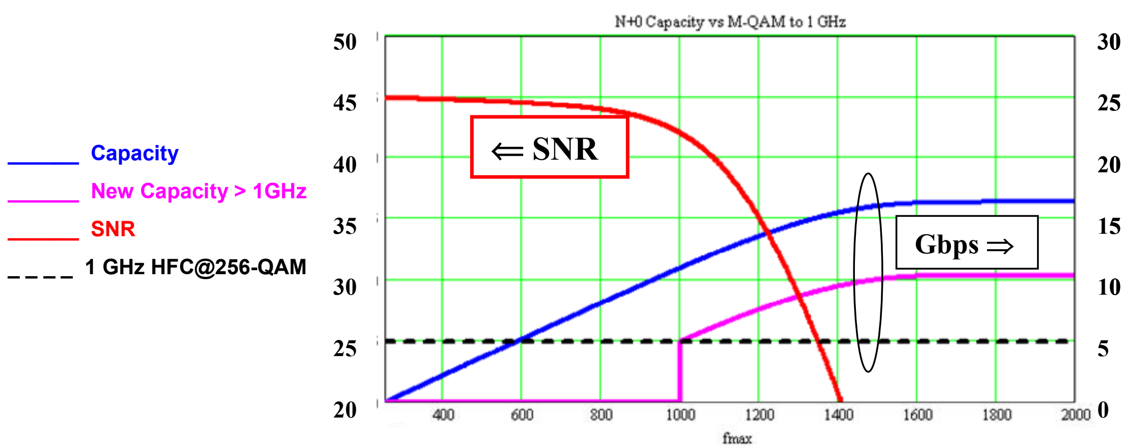

Figure 20 – Gbps: Current Use, Ideal Use, and Additional New Above 1 GHz

In Figure 20, a full forward band throughput of 256-QAM is shown, along with the theoretical capacity in Gbps (blue, right vertical axis), for a given maximum upper edge of the band shown on the x-axis. These capacities are shown along with the SNR vs. frequency delivered from a 5-tap cascade made up of taps such as that shown in Figure 19, and one coupled port from the same. The final trace (pink) recognizes the 256-QAM legacy spectrum as a given, and above that identifies new theoretical capacity potentially that can be exploited above 1 GHz in the passive segment as a function of the maximum upper frequency used. Clearly, within the first 200 MHz above 1 GHz, a Gbps of throughput can be extracted. Also apparent is how much latent capacity still exists as the cascades shrink and open up new RF bandwidth potential, considering that 256-QAM is today's maximum modulation profile. Many recent analysis have proposed use of 1024-QAM and perhaps even higher order modulations [10][11]. There are plenty of available bits per second left to be exploited in the passive infrastructure, above and beyond the current use of the bandwidth to use as "replacement" capacity should a phase 2 migration above N-Split be required.

CONCLUSION

Operators are dealing with the inevitable charge forward of traffic demand. With historical trends as a guideline, the traffic increase can be quantified and used to develop timelines and strategies for dealing with the growth. The first step of optimizing the existing 5-42 MHz is underway, with multiple operators moving ahead with DOCSIS 3.0 capabilities, including its most bandwidth-efficient modulation profile, 64-QAM. The additional tool for extracting everything possible in 5-42 MHz is S-CDMA. While it has yet to be embraced fully in North America, it is a field proven, powerful technique to make productive use of currently vacant spectrum.

Despite these 5-42 MHz optimizations, the upstream is ultimately limited by its hard cap in total bandwidth under today's diplex architecture, which highly favors downstream. Using CAGR analysis, we can estimate when this obstacle needs to be removed. A straightforward and relatively *long-term* solution, providing a decade or more of potential growth, is the 5-85 MHz N-Split architecture already called out in DOCSIS 3.0. Equipment is available now that supports N-Split, and more will become available in the very near future. Furthermore, the current generation of HFC optics using analog DFB returns is *already* capable of supporting the added bandwidth. Better yet, the new bandwidth is exceptionally clean. This fact, combined with the HFC link performance and a new generation of high fidelity CMTS receivers, makes 256-QAM usage practical with solid operating margins in the upstream. This, too, has been proven in the field using existing hardware. When applied to CAGR analysis, the results show that N-Split can capably support a decade or more of new upstream demand-based consumption. This extended timeframe positions the MSO well, offering a lengthy opportunity to evaluate technology shifts over the next ten years and plan next generation architecture steps. These benefits are derived through a low-risk transition to N-Split.

Though a long-term solution has been identified, should the next steps beyond turn out to involve simply many more years of HFC-flavored evolution, support for continued upstream traffic growth beyond the above time frame can continue by further plant segmentation or implementing a second phase of diplex adjustment that extends the band to 200 MHz or beyond, should that be necessary for larger peak service rates. If developed in a forward-looking way, new phases of diplex shifting can be done with minimal plant disruption and operational complexity. As we have shown, this bandwidth extension *also* should be capably supported, and with higher order modulation profiles such as 256-QAM, using today's generation of HFC optics and receivers on par with today's generation of low noise CMTS'. Of course, these complementary tools – new bandwidth and splitting of nodes – can and should both be considered.

Finally, if necessary, new "replacement" forward bandwidth may become easily accessible above the top end of today's forward band if the additional return imposes on downstream growth. Above the top end of the forward band, it is much simpler and more bandwidth efficient to create additional *forward* capacity than to try and push upstream signals against their architectural will in the face of many obstacles.

HFC has a long and impressive history of technology and architecture evolution, and of new services. It also has a long and undoubtedly impressive future ahead of it, capable of much more capacity exploitation, and full of potential for many new exciting services that maintain and delight customers. It is hoped that this analysis is found useful as a guideline for planning a migration strategy that fully realizes this latent potential in today's HFC networks.

ACKNOWLEDGEMENTS

The authors would like to thank Jack Moran, Robert Thompson, Charles Moore, Chris Chiu, and Jim Weineck for their invaluable contributions to many of the measured results shown in this paper.

REFERENCES

[1] Finkelstein, Jeff, "Upstream Bandwidth Futures," 2010 SCTE Cable-Tec Expo, New Orleans, LA, Oct. 20-22.

[2] Howald, Dr. Robert L., Phillip Chang, Robert Thompson, Charles Moore, Dean Stoneback, and Vipul Rathod, "Characterizing and Aligning the HFC Return Path for Successful DOCSIS 3.0 Rollouts," 2009 SCTE Cable-Tec Expo, Denver, CO, Oct 28-30.

[3] Ulm, John, Jack Moran, Daniel Howard, "Leveraging S-CDMA for Cost Efficient Upstream Capacity," SCTE Conference on Emerging Technologies, Washington DC, April 2, 2009.

[4] Howald, Dr. Robert L., "Maximizing the Upstream: The Power of S-CDMA" Communication Technology Webcast, Sept. 9, 2009.

[5] Howald, Dr. Robert L. and Michael Aviles, "Noise Power Ratio the Analytical Way," 2000 NCTA Show, New Orleans, LA.

[6] Miguelez, Phil, and Dr. Robert Howald, "Digital Due Diligence for the Upstream Toolbox," 2011 Cable Show, Chicago, IL, June 14-16.

[7] Robuck, Mike, "Cox, Motorola lay claim to new return path speed record," CedMagazine.com, March 01, 2011.

[8] Stoneback, Dean, Dr. Robert L. Howald, Joseph Glaab, Matt Waight, "Cable Modems in the Home Environment," 1998 NCTA Show, Atlanta, Ga.

[9] Dr. Robert L.,"Fueling the Coaxial Last Mile," SCTE Conference on Emerging Technologies, Washington DC, April 2, 2009.

[10] Dr. Robert L., Michael Aviles, and Amarildo Vieira, "New Megabits, Same Megahertz: Plant Evolution Dividends," 2009 Cable Show, Washington, DC, March 30-April 1.

[11] Hasse, Philipp, Dirk Jaeger and Joerg Robert, "DVB-C2 – The Second Generation Technology for Broadband Cable," 2009 Cable Show, Washington, DC, March 30-April 1.